

- 3 -

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (currently amended) A method of estimating ~~results~~ a result of a current database query, the method comprising:
 - collecting workload information related to queries that have been executed on ~~the~~ a database;
 - tracing query patterns ~~of the queries~~ in the workload information to identify a ~~usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to~~ a relative frequency with which ~~the~~ at least one given tuple with respect to at least one other given tuple was accessed ~~by queries in the workload in executing the queries;~~
 - determining a sample weight for the at least one given tuple based on ~~tuple usage for each~~ the relative access frequency of the given tuple;
 - performing a weighted sampling of the database based upon the sample weights of the at least one given tuple; and
 - executing the current database query on the at least one weighted sample to estimate ~~results~~ the result of the current database query.
2. (currently amended) The method of claim 1 wherein the weighted sampling is performed by assigning a the sample weight to ~~each~~ the at least one given tuple based on a probability of usage of the given tuple in executing the queries ~~in the workload~~.
3. (currently amended) The method of claim 2 ~~and~~ further comprising:
 - computing an aggregate over values in ~~each~~ the at least one sample tuple.
4. (currently amended) The method of claim 3 wherein the aggregate is computed by multiplying each ~~value of the values~~ by ~~the~~ an inverse of ~~the~~ a probability with which at least one corresponding ~~tuples were~~ given tuple was sampled in executing the queries.

- 4 -

5. (currently amended) The method of claim 1 wherein the sample weights ~~are~~ is a function of the a frequency of access of a the at least one given tuple and the a number of queries ~~in the workload~~ that access the given tuple in executing the queries.

6. (currently amended) The method of claim 1 wherein the relative access frequency of the at least one given tuple ~~usage~~ is stored on a page level.

7. (currently amended) A machine readable medium having instructions stored thereon, which when executed by at least one processor, cause the processor for ~~causing a machine~~ to perform a method of estimating ~~results~~ a result of a current database query, the method comprising:

collecting workload information related to queries that have been executed on ~~the a~~ a database;

tracing query patterns ~~of the queries~~ in the workload information to identify a ~~usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to~~ a relative frequency with which the at least one given tuple with respect to at least one other given tuple was accessed ~~by queries in the workload in executing the queries;~~

determining a sample weight for the at least one given tuple based on ~~tuple usage for each~~ the relative access frequency of the given tuple;

performing a weighted sampling of the database based upon the sample weights of the at least one given tuple; and

executing the current database query on the at least one ~~weighted~~ sample to estimate ~~results~~ the result of the current database query.

8. (currently amended) The machine readable medium of claim 7 wherein the sample weights ~~are~~ is a function of the a frequency of access of a the at least one given tuple and ~~the a~~ number of queries ~~in the workload~~ that access the given tuple in executing the queries.

9. (currently amended) The ~~method~~ machine readable medium of claim 7 wherein the relative access frequency of the at least one given tuple ~~usage~~ is stored on a page level.

- 5 -

10. (currently amended) A system that estimates ~~results~~ a result of a current database query, the ~~method~~ system comprising:

a module that collects workload information related to queries that have been executed on ~~the~~ a database;

a module that traces query patterns ~~of the queries~~ in the workload information to identify ~~usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to~~ a relative frequency with which ~~the~~ at least one given tuple with respect to at least one other given tuple was accessed ~~by queries in the workload in executing the queries;~~

a module that determines a sample weight for the at least one given tuple based on ~~tuple usage for each~~ the relative access frequency of the given tuple;

a module that performs a weighted sampling of the database based upon the sample weights of the at least one given tuple; and

a module that executes the current database query on the ~~weighted~~ at least one sample to estimate ~~results~~ the result of the current database query.

11-34 (canceled).

35. (currently amended) A method of estimating ~~results~~ a result of a current database query ~~and a given workload wherein the queries in the workload may have selection conditions,~~ the method comprising:

collecting workload information related to queries that have been executed on ~~the~~ a database;

tracing query patterns ~~of the queries~~ in the workload information to identify a ~~usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to~~ a relative frequency with which ~~the~~ at least one given tuple with respect to at least one other given tuple was accessed ~~by queries in the workload in executing the queries;~~

determining a sample weight for the at least one given tuple based on ~~tuple usage for each~~ the relative access frequency of the given tuple;

performing a weighted sampling of the database based upon the sample weights of the at least one given tuple;

executing the current database query on the at least one ~~weighted~~ sample to estimate ~~results~~ the result of the current database query; ~~[[,]]~~ and ~~[[,]]~~

- 6 -

generating a weighted outlier index that identifies at least one outlier tuple associated with at least one of the executed queries.

36. (currently amended) The method of claim 35 ~~and~~ wherein executing the current database query on the at least one sample to estimate the result of the current database query further comprising comprises:

calculating an aggregate of the at least one sample;

calculating an another aggregate based on the samples of the at least one outlier tuple identified in the weighted outlier index; and

combining the aggregates to estimate the result of the current database query.

37-41 (canceled).

42. (new) A method of identifying which values among a set of values to use in providing an estimated query result for a current query, the method comprising:

obtaining a usage frequency for at least one value within the set of values based on how often the at least one value was used to calculate a result in at least one executed query;

associating a weight with the at least one value that is based on the usage frequency for that value; and

selecting the at least one value whose associated weight satisfies at least one weight criterion to use in obtaining the estimated query result for the current query.

43. (new) The method of claim 42 further comprising:

calculating the estimated query result for the current query using the at least one selected value.

44. (new) The method of claim 43 wherein calculating the estimated query result for the current query using the at least one selected value further comprises:

multiplying the at least one selected value by an inverse of a probability that the selected value was used to calculate the result in the at least one executed query.

- 7 -

45. (new) The method of claim 42 wherein obtaining a usage frequency for at least one value within the set of values based on how often the at least one value was used to calculate a result in at least one executed query further comprises:

obtaining the usage frequency for the at least one value based on at least one of either how many times the value was used to calculate the result in the at least one executed query, how many of the at least one executed query used the value to calculate the result, and a probability that the value was used to calculate the result in the at least one executed query.

46. (new) The method of claim 42 further comprising:

generating an outlier index that identifies at least one outlier value associated with the at least one executed query.

47. (new) The method of claim 46 wherein executing the current database query on the at least one sample to estimate the result of the current database query further comprises:

calculating the estimated query result for the current query using the at least one selected value and the at least one outlier value identified in the outlier index.

48. (new) The method of claim 46 wherein generating an outlier index that identifies at least one outlier value associated with the at least one executed query further comprises:

generating the outlier index to identify the at least one outlier value associated with the at least one executed query that is executed more frequently than at least one other executed query.